Lie-Access Neural Turing Machines

Greg Yang and Alexander M. Rush
{gyang@college,srush@seas}.harvard.edu
Harvard University
Cambridge, MA 02138, USA

arXiv:1611.02854v1 [cs.NE] 9 Nov 2016

Abstract

Recent work has demonstrated the effectiveness of employing explicit external memory structures in conjunction with deep neural models for algorithmic learning (Graves et al., 2014; Weston et al., 2014). These models utilize differentiable versions of traditional discrete memory-access structures (random access, stacks, tapes) to provide the variable-length storage necessary for computational tasks. In this work, we propose an alternative model, Lie-access memory, that is explicitly designed for the neural setting. In this paradigm, memory is accessed using a continuous head in a key-space manifold. The head is moved via Lie group actions, such as shifts or rotations, generated by a controller, and soft memory access is performed by considering the distance to keys associated with each memory. We argue that Lie groups provide a natural generalization of discrete memory structures, such as Turing machines, as they provide inverse and identity operators while maintain differentiability. To experiment with this approach, we implement several simplified Lie-access neural Turing machine (LANTM) with different Lie groups. We find that this approach is able to perform well on a range of algorithmic tasks.

1 Introduction

Pioneering work on neural Turing machines (NTMs) (Graves et al., 2014; 2016) and memory networks (MemNNs) (Weston et al., 2014) has demonstrated that explicit external memories can be effectively utilized within deep neural networks, without sacrificing the end-to-end nature of these models. Recently researchers have studied many variations of external neural memories. Early work has explored tasks such as question answering (Weston et al., 2014; Sukhbaatar et al., 2015; Kumar et al., 2015), algorithm learning (Graves et al., 2014; Kalchbrenner et al., 2015; Kaiser & Sutskever, 2015; Kurach et al., 2015; Zaremba & Sutskever, 2015; Grefenstette et al., 2015; Joulin & Mikolov, 2015), machine translation (Kalchbrenner et al., 2015), and others. This work has the potential to extend deep networks beyond the limitations of fixed-length encodings such as standard recurrent neural networks (RNNs).

A shared theme in many of these works (and earlier exploration of neural memory) is to re-frame traditional memory access paradigms to be continuous and possibly differentiable to allow for simple backpropagation. In MemNNs, traditional *random-access* memory is replaced with a ranking approach that finds the most likely memory. In the work of Grefenstette et al. (2015), classical *stack-based* memory is replaced by a soft differentiable stack data-structure. In NTMs, *sequential* local-access memory is simulated by an explicit tape data structure.

We propose to instead specify a class of memory access techniques specifically designed for the neural context, that do not attempt to directly model known discrete access methods. In particular, inspired by discrete Turing machines, we argue that a neural memory should maintain the following properties: (A) maintaining differentiability for end-to-end training and (B) provide a robust relative indexing scheme (perhaps in addition to random-access).

Lie groups, mathematical groups with differentiable group operations, provide a natural structure for neural memory access. By definition, their differentiability satisfies the concerns of Criterion

A. Additionally the group axioms provide identity, invertibility, and associativity, all of which are desirable properties for a relative indexing scheme, and all of which are satisfied by standard Turing machines. Notably though, simple group properties like invertibility are not satisfied by neural Turing machines, differentiable neural computers, or even by simple soft-tape machines.

To experiment with this approach, we implement a neural Turing machine with an LSTM controller and several versions of Lie-access memory, which we call Lie-access neural Turing machines (LANTM). The details of these models are exhibited in Section 4. Our main experimental results are presented in Section 5. The LANTM model is able to learn non-trivial algorithmic tasks such as copying and permutating sequences with a similar accuracy of more traditional memory-based approaches, and significantly better than fixed memory LSTM models. The memory structures and key transformation learned by the model resemble interesting continuous space representations of traditional discrete memory data structures.

2 Background: Recurrent Neural Networks with Memory

We focus particularly on the case of an RNN controller of an abstract memory network. Formally, an RNN is a differentiable function $\text{RNN} : \mathcal{X} \times \mathcal{H} \to \mathcal{H}$, where $\mathcal{X}$ is an arbitrary input space and $\mathcal{H}$ the space of internal states. On input $(x^{(1)}, \ldots, x^{(T)}) \in \mathcal{X}^T$ and with initial state $h^{(0)} \in \mathcal{H}$, the RNN transitions into states $h^{(1)}, \ldots, h^{(T)}$,

$$h^{(t)} := \text{RNN}(x^{(t)}, h^{(t-1)}).$$

These states can then be used for downstream tasks, e.g. sequence prediction $(y^{(1)}, \ldots, y^{(T)})$, based on an additional transformation and prediction layer $y^{(t)} = F(h^{(t)})$ such as a linear-layer followed by a softmax. RNNs can be trained end-to-end by backpropagation-through-time (BPTT) (Werbos, 1990). In practice, we use a particular variant of RNN called the long short-term memory (LSTM) network (Hochreiter & Schmidhuber, 1997). LSTM's hidden state consists of two variables $(c^{(t)}, h^{(t)})$, where $h^{(t)}$ is also the output to the external world; we however use the above notation for simplicity.

In models with external memories, the RNN can serve as the controller (Graves et al., 2014; Grefenstette et al., 2015; Zaremba & Sutskever, 2015). This means that (1) the entire system carries state over time from both the RNN and the external memory, and (2) the RNN controller collects readings from and computes additional instructions to the external memory. The system is now described by the recurrence:

$$\begin{aligned} h^{(t)} &:= \text{RNN}([x^{(t)}; \rho^{(t-1)}], h^{(t-1)}), \\ \Sigma^{(t)}, \rho^{(t)} &:= \text{RW}(\Sigma^{(t-1)}, h^{(t)}), \end{aligned}$$

where Σ is the memory state, and $\rho^{(t)}$ is the memory being read. As above, output prediction can be made using a standard final layer $y^{(t)} := F(h^{(t)})$. Writing occurs in the mutation of Σ at each time step. Throughout this work, Σ will take the form $\{\ldots (k_i, v_i, s_i) \ldots\}$ where $v_i \in \mathbb{R}^m$ are the memories, k_i is an arbitrary key, and $s_i \in \mathbb{R}$ is a memory strength, described below.

In order for the model to be trainable, the memory function RW must also be differentiable. Several forms of differentiable memory have been proposed in the literature. We begin by describing two simple forms: (neural) random-access memory and (neural) tape-based memory. For this section, we focus on the read step and assume Σ is fixed.

Random-Access Memory Random-access memory consists of using a now standard attention-mechanism or soft MemNN to read a memory based on a controller. This description is similar to the work of Miller et al. (2016). The controller hidden state is used to output a random-access pointer $\tilde{q}(h^{(t)})$ that determines a weighting of memory vectors via dot products with the corresponding keys. This weighting in turn determines the read values via convex combination (attention).

$$w_i := \frac{s_i \exp \langle \tilde{q}, k_i \rangle}{\sum_j s_j \exp \langle \tilde{q}, k_j \rangle} \quad \text{for all } i, \qquad\qquad \rho := \sum_i w_i v_i$$

The final read memory can be seen as determined by how "close" the read pointer was to each of the keys.

Tape-Based Memory Neural memories can also keep around a state indicating a current read position. Following notation from Turing machines, we call this state the *head*, q. In the simplest case the recurrence now has the form,

$$\Sigma', q', \rho = \mathrm{RW}(\Sigma, q, h),$$

although we might incorporate multiple heads.

In the simplest case of soft tape-based memory (a naive version of the much more complicated neural Turing machine), the keys k_i indicate positions along a tape with $k_i = i$. The head q is a probability distribution over locations of the memory i. It determines the read value by directly specifying the weight, i.e. $w_i = q_i$. The controller can only "shift" the head by outputting a kernel $K = K(h) = (K_{-1}, K_0, K_{+1})$ in the probability simplex Δ^2 and applying convolution.

$$\begin{aligned} q' &:= q * K, \qquad \text{i.e.} & q'_j &= q_{j-1}K_{+1} + q_j K_0 + q_{j+1}K_{-1}, \\ w_i &:= q'_i \ \text{ for all } i, & \rho &:= \sum_i w_i v_i. \end{aligned}$$

We can view this as the soft version of a single-step discrete Turing machine where the kernel can softly shift the "head" of the machine one to the left, one to the right, or remain in the same location.

3 Lie Groups for Memory

Let us revisit now the classical discrete Turing machine and the movement of its head over a tape. In particular, between reads, the head can move any integer number of steps left or right. Moving $a + b$ steps and then c steps eventually puts the head at the same location as moving a steps and then $b + c$ steps — i.e. the head movement is associative. In addition, the machine should be able to reverse a head shift, for example, in a stack simulation algorithm, going from push to pop — i.e. each head movement should also have a corresponding inverse. Finally, the head should also be allowed to stay put, for example, to read a single data item and use it for multiple time points.

These movements correspond directly to *group actions*: the possible head movements should be associative, and contain inverse and identity elements. This group acts on the set of possible head locations. In a TM, the set of $\mathbb{Z}$-valued head movement acts on the set of locations on the $\mathbb{Z}$-indexed infinite tape. By our reasoning above, if a Turing machine is to store data contents at points in a general space $\mathcal{K}$ (instead of an infinite $\mathbb{Z}$-indexed tape), then its head movements should form a group and act on $\mathcal{K}$ via group actions.

For a neural memory system, we desire most things to be differentiable (or almost everywhere differentiable). The notion of "differentiable groups" is well-studied in mathematics, where they are known as *Lie groups*, and "differentiable group actions" are correspondingly called *Lie group actions*. In our case, using Lie group actions as generalized head movements on a general space (more accurately, manifolds) would most importantly mean that we can take derivatives of these movements and perform the usual backpropagation algorithm.

4 Lie-Access Neural Turing Machines

These properties motivate us to propose, Lie access as an alternative formalism to popular soft-versions of neural memory systems, such as probabilistic tapes, which surprisingly do not satisfy invertibility and often do not provide an identity.[1] Our Lie-access memory will consist of a set of

[1]The Markov kernel convolutional soft head shift mechanism proposed in Graves et al. (2014) and explained in section 2 does not in general have inverses. Indeed, the authors reported problems with the soft head losing "sharpness" over time, which they dealt with by "sharpening coefficients." In the followup

points in a manifold $\mathcal{K}$. We replace the discrete head, with continuous head, $q \in \mathcal{K}$. The head will move based on a set of Lie group actions $a \in \mathcal{A}$ generated by the controller. To read memories, we will rely on a distance measure in this space, $d : \mathcal{K} \times \mathcal{K} \to \mathbb{R}^{\geq 0}$.[2] Together these properties describe a general class of possible neural memory architectures.

Formally a Lie-access neural Turing machine (LANTM) computes the following function,

$$\Sigma', q', q'_{(w)}, \rho := \mathrm{RW}(\Sigma, q, q_{(w)}, h)$$

where $q, q_{(w)} \in \mathcal{K}$ are resp. read and write heads, and Σ is the memory itself. We implement Σ as above as a weighted dictionary $\Sigma = \{\ldots, (k_i, v_i, s_i), \ldots\}$ where $k_i \in \mathcal{K}$ is a key, $v_i \in \mathbb{R}^m$ is the value, and $s_i \in [0, 1]$ is a strength value, such that when $s_i = 0$, the entry (k_i, v_i) can be ignored.

4.1 Addressing Procedure

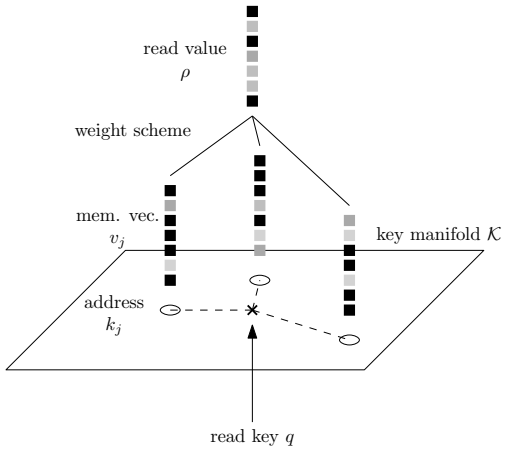

Figure 1: Retrieval of value from memory via a key. Weightings with unit sum are assigned to different memories depending on the distances from the addresses to the read key. The weighted arithmetic mean is emitted as the final read value. Both InvNorm and softmax schemes follow this method, but differ in computing the weightings.

The LANTM maintains a read head q which at every step is updated to q' and is used to read from the memory table. This update occurs by a Lie-access procedure where an element from the Lie action group $\mathcal{A}$ acts smoothly on the key space $\mathcal{K}$.

We utilize parameterized action transformation, $A : \mathcal{H} \mapsto \mathcal{A}$, to produce the Lie action, $A(h) \in \mathcal{A}$. In the simplest case, the head is then updated based on this action (here $\cdot$ denotes group action):

$$q' := A(h) \cdot q$$

For instance, consider two possible Lie groups:

- A shift group $\mathbb{R}^2$ acting additively on $\mathbb{R}^2$. This means that $\mathcal{A} = \mathbb{R}^2$ so that $A(h) = (\alpha, \beta)$ acts upon a key $q = (x, y)$ by
$$A(h) \cdot q = (\alpha, \beta) + (x, y) = (x + \alpha, y + \beta).$$
- The rotation group $SO(3)$ acting on the sphere $S^2 = \{v \in \mathbb{R}^3 : \|v\| = 1\}$. Each rotation can be described by its axis ξ and angle θ. An action $(\xi, \theta) \cdot q$ is just the appropriate rotation of the point q, and is given by Rodrigues' rotation formula

$$A(h) \cdot q = (\xi, \theta) \cdot q = q \cos\theta + (k \times q) \sin\theta + k\langle k, q\rangle(1 - \cos\theta).$$

4.2 Reading Memories

Recall that memories are stored in Σ, each with a key, k_i, memory vector, v_i, and strength, s_i. Memories are read using a soft weighting scheme based on their key address. The weighting is defined

work, Graves et al. (2016) utilize a *temporal memory link matrix* for actions. They note, "the operation Lw smoothly shifts the focus forwards to the locations written ... whereas $L^\top w$ shifts the focus backwards" but do not enforce this as a true inverse. They also explicitly do not include an identity, noting "Self-links are excluded (the diagonal of the link matrix is always 0)".

[2] This metric should satisfy a compatibility relation with the Lie group action When points $x, y \in X$ are simultaneously moved by the same Lie group action v, their distance should stay the same (One possible mathematical formalization is that X should be Riemannian manifold and the Lie group should be a subgroup of X's isometry group.): $d(vx, vy) = d(x, y)$. This condition ensures that if the machine writes a sequence of data along a "straight line" at points $x, vx, v^2x, \ldots, v^kx$, then it can read the same sequence by emitting a read location y close to x and then follow the "v-trail" $y, vy, v^2y, \ldots, v^ky$.

by a weight vector w, and the read vector is a linear combination over memories: $\rho' := \sum_i w_i v_i$. While there are many possible weighting schemes, we use one based on the distance of each memory from the head in key-space assuming a metric d on $\mathcal{K}$. One method is to compute weight w of each memory using the polynomial law for all i as,

$$w_i := \frac{s_i d(q', k_i)^{-2}}{\sum_j s_j d(q', k_j)^{-2}}$$

with the convention that it takes the limit value when $q \to k_i$ for some i. We call this method of converting a read key to a set of weighting *InvNorm* for short.

An alternative is to use the more common exponential law/annealed softmax weight scheme, which computes the weights w_i as

$$w_i := \frac{s_i \exp(-d(q', k_i)^2/T)}{\sum_j s_j \exp(-d(q', k_j)^2/T)},$$

where T is a *temperature* emitted by the controller at time t that represents the certainty of its reading, i.e. higher T creates more uniform w.

4.3 Writing Memories

The writing procedure is similar to reading. The LANTM maintains a separate *write head* $q_{(w)}$ that moves analogously to the read head, i.e. with action function $a_{(w)}$ and updated value $q'_{(w)}$. At each call to RW, a new memory is automatically appended to Σ with $k = q'_{(w)}$. The corresponding v and s are created by MLP's $v(h) \in \mathbb{R}^m$ and $s(h) \in [0, 1]$ taking h as input. After writing, the new memory set is,

$$\Sigma' := \Sigma \cup \{(q'_{(w)}, v(h), s(h))\}.$$

No explicit erase mechanism is provided, but to erase a memory (k, v, s), the controller may just write $(k, -v, s)$.

4.4 Interpolations

In practice we can combine this relative addressing procedure with direct random-access by outputting an interpolation. Write $t(h) \in [0, 1]$ for the interpolation gate and $\tilde{q}(h) \in \mathcal{K}$ for our proposed random-access layer. For key space manifolds $\mathcal{K}$ like $\mathbb{R}^n$, there's a well defined straight-line interpolation between two points, so we can set

$$q' := A \cdot (t \times q + (1 - t) \times \tilde{q})$$

where we have omitted the implied dependence on h. For other manifolds like the spheres S^n that have well-behaved projection functions $\pi : \mathbb{R}^n \to S^n$, we can just project the straight-line interpolation to the sphere:

$$q' := A \cdot \pi(t \times q + (1 - t) \times \tilde{q}).$$

In the case of a sphere S^n, π is just L_2-normalization.[3]

Additionally, we also extend the model to interpolate the last action $A^{(t-1)}$ with a candidate action $A(h)$ via a gate $r(h) \in [0, 1]$ to produce the final action $A^{(t)}$ (replacing A above). This is so that the controller may move a head in "a straight line" by just saturating both t and r. For example, for the translation group $\mathbb{R}^n$, we have straight-line interpolation as above, $A^{(t)} := r \times A + (1 - r) \times A^{(t-1)}$.

For the rotation group $SO(n)$, each rotation is represented by its axis $\xi \in S^{n-1}$ and angle $\theta \in (-\pi, \pi]$, and we just interpolate each separately $\xi^{(t)} := \pi(r \times \xi + (1-r) \times \xi^{(t-1)})$ and $\theta^{(t)} := r \times \theta + (1-r) \times \theta^{(t-1)}$. where π is L_2-normalization.[4]

[3] Technically, in the sphere case, $\operatorname{dom} \pi = \mathbb{R}^d - \{0\}$. But in practice one almost never gets 0 from a straight-line interpolation, so computationally this makes little difference.

[4] There is, in fact, a canonical way to interpolate the most common Lie groups, including all of the groups mentioned above, based on the exponential map and the Baker-Campbell-Hausdorff formula (Lee,

5 Experiments

We experiment with this new approach to memory on a set of algorithmic learning tasks. We are particularly interested in: (a) how Lie-access memory can trained, (b) whether it can learn be effectively utilized for algorithmic learning, and (c) what internal structures the model learns compared to systems based directly on discrete memory. In particular Lie access is not equipped with an explicit stack or tape, so it needs to learn continuous patterns that capture these properties.

Setup. Our experiments utilize an LSTM controller in a version of the encoder-decoder setup (Sutskever et al., 2014), i.e. an encoding input pass followed by decoding output pass. The encoder reads and writes memories at each step; the decoder only reads memories. The encoder is given $\langle s \rangle$, followed by an the input sequence, and then $\langle / s \rangle$ to terminate input. In the decoder, the system is not re-fed its output or the correct symbol, i.e. we do not use teacher forcing. The decoder must correctly emit an end-of-output symbol $\langle / e \rangle$ to terminate.

Models and Baselines. We implement three main baseline models including: (a) a standard *LSTM* encoder-decoder, without explicit external memory, (b) a random access memory network, *RAM* using the key-value formulation as described the background, roughly analogous to an attention-based encoder-decoder, (c) an interpolation of a *RAM/Tape*-based memory network as described in the background, i.e. a highly simplified version of a true NTM. Our models include four versions of Lie-access memory. The main model, *LANTM*, has an LSTM controller, with a shift group $\mathcal{A} = \mathbb{R}^2$ acting additively on $\mathcal{K} = \mathbb{R}^2$. We also consider a model *SLANTM* with spherical memory, utilizing a rotation group $\mathcal{A} = SO(3)$ acting on keys in the sphere $\mathcal{K} = S^2$. For both of the models, we experiment with indexing using *InvNorm* (default) and with an annealed *softmax*.[5]

Model Setup. For all tasks, the LSTM baseline has 1 to 4 layers, each with 256 cells. Each of the other models has a single-layer, 50-cell LSTM controller, with memory width (i.e. the size of each memory vector) 20. Other parameters such as learning rate, decay, and intialization are found through grid search. Further hyperparameter details are give in the appendix.

Tasks. Our experiments are on a series of algorithm tasks shown in Table 1a. The tasks copy, reverse, and bigram flip tasks are based on Grefenstette et al. (2015); tasks double and interleaved add are designed in a similar vein. Additionally we also include two harder tasks repeat copy and priority sort. In repeat copy each model must repeat a sequence of length 20, N times. In priority sort, each item of the input sequence is given a priority, and the model must output them in priority order. The models are trained with random 32K examples and tested on 3.2K random examples.

Prediction is performed via argmax at each step; training is done by minimizing negative log likelihood with RMSProp. For each task, we train each model with the same number of samples. In particular it is not the case that we train till convergence as was done in e.g. Grefenstette et al. (2015). To evaluate the performance of the models, we compute the fraction of characters correctly predicted and the fraction of all answers completely correctly predicted, respectively called fine/ coarse following Grefenstette et al. (2015). We score the models on randomly generated in-sample 1x data and out-of-sample 2x data, with sequence lengths $2k$ (or repeat number $2N$ in the case of repeat copy) to test the generalization of the system.

Results. Main results comparing the different memory systems and read computations on a series of tasks are shown in Table 1b. Consistent with previous work the fixed-memory LSTM system is able to do reasonably well on the 1x samples, but fails consistently when required to generalize to the 2x samples, unable to solve any 2x problem correctly. The RAM (attention-based) and the

2012), but the details are outside the scope of this paper and the computational cost, while acceptable in control theory settings, is too hefty for us. Interested readers are referred to Shingel (2009) and Marthinsen (1999).

[5] Note that the read weight calculation of a SLANTM with softmax is essentially the same as the RAM model: For head q', $\exp(-d(q', k_i)^2/T) = \exp(-\|q' - k_i\|^2/T) = \exp(-(2 - 2\langle q', k_i \rangle)/T)$, where the last equality comes from $\|q'\| = \|k_i\| = 1$. Therefore the weights $w_i = \frac{s_i \exp(-d(q',k_i)^2/T)}{\sum_j s_j \exp(-d(q',k_j)^2/T)} = \frac{s_i \exp(-2\langle q',k_i \rangle/T)}{\sum_j s_j \exp(-2\langle q',k_j \rangle/T)}$, which is the RAM weighting scheme.

Task	Input	Output	Size k	$\lvert\mathcal{V}\rvert$
1 - Copy	$a_1 a_2 a_3 \cdots a_k$	$a_1 a_2 a_3 \cdots a_k$	$[2, 64]$	128
2 - Reverse	$a_1 a_2 a_3 \cdots a_k$	$a_k a_{k-1} a_{k-2} \cdots a_1$	$[2, 64]$	128
3 - Bigram Flip	$a_1 a_2 a_3 a_4 \cdots a_{2k-1} a_{2k}$	$a_2 a_1 a_4 a_3 \cdots a_{2k} a_{2k-1}$	$[1, 16]$	128
4 - Double	$a_1 a_2 \cdots a_k$	$2 \times \lvert a_k \cdots a_1 \rvert$	$[2, 40]$	10
5 - Interleaved Add	$a_1 a_2 a_3 a_4 \cdots a_{2k-1} a_{2k}$	$\lvert a_{2k} a_{2k-2} \cdots a_2 \rvert + \lvert a_{2k-1} a_{2k-3} \cdots a_1 \rvert$	$[2, 16]$	10
6 - Odd First	$a_1 a_2 a_3 a_4 \cdots a_{2k-1} a_{2k}$	$a_1 a_3 \cdots a_{2k-1} a_2 a_4 \cdots a_{2k}$	$[1, 16]$	128
7 - Repeat Copy	$\overline{N} a_1 \cdots a_{20}$	$a_1 \cdots a_{20} \cdots a_1 \cdots a_{20}$ (N times)	$N \in [1, 5]$	128
8 - Priority Sort	$\overline{5} a_5 \overline{2} a_2 \overline{9} a_9 \cdots$	$a_1 a_2 a_3 \cdots a_k$	$[2, 10]$	128

(a) Task descriptions and parameters. $\lvert a_k \cdots a_1 \rvert$ means the decimal number repesented by decimal digits $a_k \cdots a_1$. Arithmetic tasks have all numbers formatted with the least significant digits *on the left* and with zero padding. The *Double* task takes an integer $x \in [0, 10^k)$ padded to k digits and outputs $2x$ in $k+1$ digits, zero padded to $k+1$ digits. The *Interleaved Add* task takes two integers $x, y \in [0, 10^k)$ padded to k digits and interleaved, forming a length $2k$ input sequence and outputs $x+y$ zero padded to $k+1$ digits. The last two tasks use numbers in unary format: $\overline{N}$ is the shorthand for a length N sequence of a special symbol @, encoding N in unary, e.g. $\overline{3} =$ @@@.

	Base		Memory				Lie							
	LSTM		RAM		RAM/Tape		LANTM		LANTM-s		SLANTM		SLANTM-s	
	1x	2x	1x	2x	1x	2x	1x	2x	1x	2x	1x	2x	1x	2x
1	74/50	22/0	⋆	60/0	99/97	40/0	⋆	★	⋆	★	⋆	★	⋆	99/94
2	94/77	30/0	⋆	53/3	⋆	53/0	⋆	**99/99**	⋆	80/55	⋆	93/45	⋆	96/55
3	86/80	28/0	⋆	71/9	⋆	65/5	⋆	★	⋆	80/55	⋆	82/19	⋆	83/24
4	98/85	43/0	⋆	72/9	⋆	69/1	⋆	★	⋆	★	⋆	94/73	⋆	99/92
5	99/98	56/0	⋆	58/7	⋆	77/19	⋆	**99/99**	58/31	25/0	⋆	**99/99**	⋆	99/93
6	85/67	29/0	⋆	40/4	⋆	64/13	⋆	66/**33**	99/85	**87**/18	⋆	45/4	97/60	84/12
7	54/31	29/0	⋆	99/10	⋆	95/40	84/50	48/0	99/92	98/5	99/84	98/15	⋆	**99/77**
8	⋆	43/0	⋆	73/11	⋆	66/18	⋆	97/72	⋆	97/74	⋆	93/60	⋆	**97/76**

(b) Main results. Numbers represent the accuracy percentages on the fine/coarse evaluations of the tasks. Symbol ⋆ indicates exact 100% accuracy (Fine scores above 99.5 are not rounded up). Baselines are described in the paper. 1x scores are obtained from test sequences of same lengths k (or N in repeat copy) as used in training. 2x scores are tested on sequence lengths $2k$ (or $2N$ in repeat copy) for generalization. For example, copy task 2x testing involves lengths 65 to 128. LANTM and SLANTM use InvNorm while LANTM-s and SLANTM-s use softmax weighting scheme. Best 2x fine and coarse scores are bolded.

RAM/tape hybrid are much stronger baselines, succeeding completely on 1x samples, but performing less well on the 2x generalization. With regard to 2x generalization, RAM does slightly better than RAM/tape on copy, reverse, bigram Flip, and double tasks, but RAM/tape does much better than RAM on the rest — interleaved add, odd first, repeat copy, and priority sort — which can be seen as more complex tasks than the previous cohort.

The last four columns illustrate the performance of the LANTM models. We found the InvNorm LANTM model to be the most effective, achieving near perfect accuracy on the first five tasks. InvNorm seems to give improvement here over the more common softmax approach to memory access. We also compare results with the sperical SLANTM model. A spherical memory system is a bit atypical, but these results demonstrate that different Lie groups can easily be substituted in, and trained. Interestingly for SLANTM, the softmax achieves better generalization. SLANTM is particularly effective on task 7 (repeat copy), getting the best generalization accuracies.

Qualitative. We did further qualitative analysis of the different Lie-access techniques to see how the models were learning the underlying tasks. Figure 2 shows two diagrams of the LANTM model of the tasks of priority sort and repeat copy. Figure 3 shows two diagrams of the SLANTM model for the same two tasks. They demonstrate that the models indeed learn interesting algorithms.

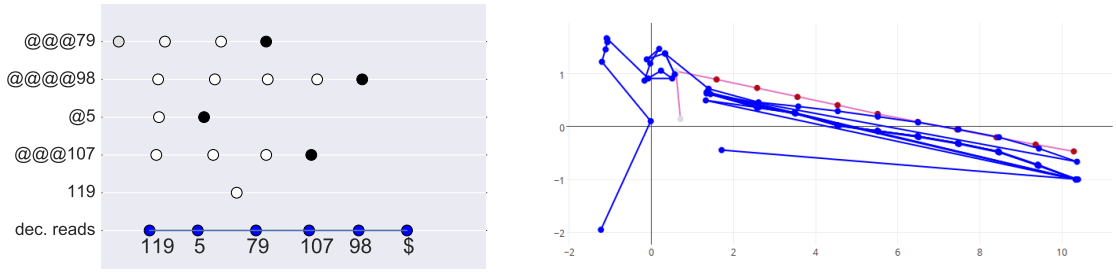

Figure 2: Analysis of the LANTM model. (a) PCA projection to 1D for the memories Σ and read heads $q^{(t)}$ of LANTM for the priority sort task. Recall in priority sort the encoder reads a priority and then a value; the decoder must output these values in priority order. In this example the sequence is $[\overline{2}, 79, \overline{3}, 98, \overline{1}, 5, \overline{3}, 107, \overline{0}, 119]$. The last line indicates read head movement. All other lines from top to bottom are the write patterns when the model receives the input subsequence $\overline{p}a_p$. Fill indicates the strength s_i. Position indicate the projected key k_i. The controller remembers the item 119 instead of writing it to Lie-access memory. (b) Raw writes (red) and reads (blue) from LANTM on repeat copy. Red lines indicate the writes occurring along a straight line during the encoding phase. Blue lines indicate the zipping back and forth of read heads in the process of copying the input sequence 6 times.

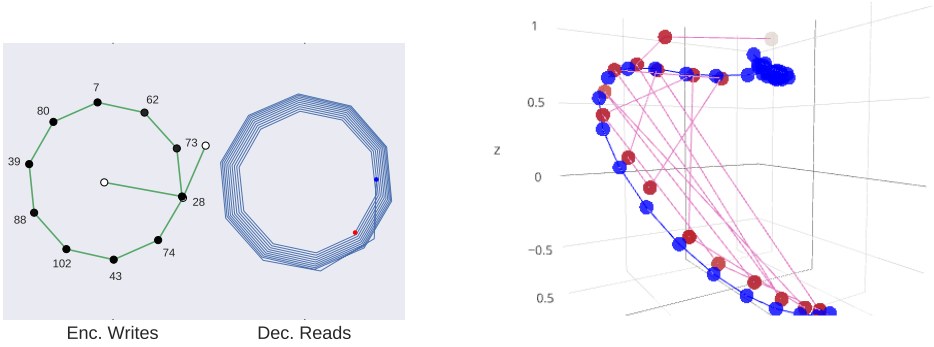

Figure 3: Analysis of the SLANTM model. (a) PCA projection to 2D, from the sphere, of the memories Σ and read heads $q^{(t)}$ of SLANTM for the task of repeat copy. Task is to repeatedly output the sequence 10 times. Input is 10 repetitions of special symbol @ followed by [28, 74, 43, 102, 88, 39, ...]. Left: the writes during the encoder. Fill indicates strength s_i, number indicates the input value read. Right: reads during the decoder phase. Starting from the blue dot, the reads move clockwise around the sphere, and end at the red dot. The model implements a cyclic list data structure. (b) Raw writes (red) and reads (blue) from SLANTM on a non-unary encoded variant of the priority sort task. Red lines indicate the random-access movements of the write-head $q_{(w)}$ to place points along a sub-manifold of $\mathcal{K}$ during the encoding phase. Blue lines indicate the simple traversal of this space during decoding.

6 Conclusion

This paper introduces Lie-access memory as an alternative neural memory access paradigm, and explored several different implementations of this approach. LANTMs follow similar axioms as discrete Turing machines while providing differentiability. Experiments show that simple models can learn algorithmic tasks. Internally these models naturally learn equivalence of standard data structures like stack and cyclic lists. In future work we hope to experiment with more groups and to scale these methods to more difficult reasoning tasks.

References

Alex Graves, Greg Wayne, and Ivo Danihelka. Neural Turing Machines. *arXiv:1410.5401 [cs]*, October 2014. URL http://arxiv.org/abs/1410.5401. arXiv: 1410.5401.

Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538 (7626):471–476, 2016.

Edward Grefenstette, Karl Moritz Hermann, Mustafa Suleyman, and Phil Blunsom. Learning to Transduce with Unbounded Memory. *arXiv:1506.02516 [cs]*, June 2015. URL http://arxiv.org/abs/1506.02516. arXiv: 1506.02516.

Sepp Hochreiter and Jrgen Schmidhuber. Long Short-Term Memory. *Neural Comput.*, 9(8):1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL http://dx.doi.org/10.1162/neco.1997.9.8.1735.

Armand Joulin and Tomas Mikolov. Inferring Algorithmic Patterns with Stack-Augmented Recurrent Nets. *arXiv:1503.01007 [cs]*, March 2015. URL http://arxiv.org/abs/1503.01007. arXiv: 1503.01007.

ukasz Kaiser and Ilya Sutskever. Neural GPUs Learn Algorithms. *arXiv:1511.08228 [cs]*, November 2015. URL http://arxiv.org/abs/1511.08228. arXiv: 1511.08228.

Nal Kalchbrenner, Ivo Danihelka, and Alex Graves. Grid Long Short-Term Memory. *arXiv:1507.01526 [cs]*, July 2015. URL http://arxiv.org/abs/1507.01526. arXiv: 1507.01526.

Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, and Richard Socher. Ask Me Anything: Dynamic Memory Networks for Natural Language Processing. *arXiv:1506.07285 [cs]*, June 2015. URL http://arxiv.org/abs/1506.07285. arXiv: 1506.07285.

Karol Kurach, Marcin Andrychowicz, and Ilya Sutskever. Neural Random-Access Machines. *arXiv:1511.06392 [cs]*, November 2015. URL http://arxiv.org/abs/1511.06392. arXiv: 1511.06392.

John Lee. *Introduction to Smooth Manifolds*. Number 218 in Graduate Texts in Mathematics. Springer, 2 edition, 2012. ISBN 978-1-4419-9981-8.

A. Marthinsen. Interpolation in Lie Groups. *SIAM Journal on Numerical Analysis*, 37(1):269–285, January 1999. ISSN 0036-1429. doi: 10.1137/S0036142998338861. URL http://epubs.siam.org/doi/abs/10.1137/S0036142998338861.

Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. Key-value memory networks for directly reading documents. *CoRR*, abs/1606.03126, 2016. URL http://arxiv.org/abs/1606.03126.

Tatiana Shingel. Interpolation in special orthogonal groups. *IMA Journal of Numerical Analysis*, 29(3):731–745, July 2009. ISSN 0272-4979, 1464-3642. doi: 10.1093/imanum/drn033. URL http://imajna.oxfordjournals.org/content/29/3/731.

Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. End-To-End Memory Networks. *arXiv:1503.08895 [cs]*, March 2015. URL http://arxiv.org/abs/1503.08895. arXiv: 1503.08895.

Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to Sequence Learning with Neural Networks. *arXiv:1409.3215 [cs]*, September 2014. URL http://arxiv.org/abs/1409.3215. arXiv: 1409.3215.

Paul J. Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=58337.

Jason Weston, Sumit Chopra, and Antoine Bordes. Memory Networks. *arXiv:1410.3916 [cs, stat]*, October 2014. URL http://arxiv.org/abs/1410.3916. arXiv: 1410.3916.

Wojciech Zaremba and Ilya Sutskever. Reinforcement Learning Neural Turing Machines - Revised. *arXiv:1505.00521 [cs]*, May 2015. URL http://arxiv.org/abs/1505.00521. arXiv: 1505.00521.

Appendices

A Experimental details

We obtain our results by performing a grid search over the grid specified in Table A.1 and also over seeds 1 to 3, and take the best scores. We bound the norm of the LANTM head shifts by 1, whereas we try both bounding and not bounding the angle of rotation in our grid for SLANTM. We initialize the Lie access models to favor Lie access over random access through the interpolation mechanism discussed in section 4.4.

The RAM model read mechanism is as discussed in section 2, and writing is done by appending new (k, v, s) tuples to the memory Σ. The only addition to this model in RAM/tape is that left and right keys are now computed using shifted convolution with the read weights:

$$k_L := \sum_i w_{i+1} k_i$$

$$k_R := \sum_i w_{i-1} k_R$$

and these keys k_L and k_R are available (along with the random access key output by the controller) to the controller on the next turn to select from via interpolation.

	rnn size	embed	decay delay	init	learning rate	key dim	angle bound
LANTM(-s)	50×1	14	{300, 600}	{1, *}	{0.01, 0.02, 0.04}	2	-
SLANTM(-s)	50×1	14	{300, 600}	{1, *}	{0.01, 0.02, 0.04}	3	{true, false}
RAM(/tape)	50×1	14	{300, 600}	{1, *}	{0.01, 0.02, 0.04}	{2, 20}	-
LSTM	$256 \times$ {1 to 4}	128	{500, 700}	*	$2 \times 10^{-\{1 \text{ to } 4\}}$	-	-

Table A.1: Parameter grid for grid search. LANTM(-s) means LANTM with invnorm or SoftMax; similarly for SLANTM(-s). RAM(/tape) means the ram and hybrid ram/tape models. Initialization: both initialization options set the forget gate of the LSTMs to 1. The number 1 in the init column means initialization of all other parameters uniformly from $[-1, 1]$. The symbol * in init column means initialization of all linear layers were done using the torch default, which initializes weights uniformly from $(-\kappa, \kappa)$, where κ is (input size)$^{-1/2}$. For models with memory, this means that the LSTM input to hidden layer is initialized approximately from $[-0.07, 0.07]$ (other than forget gate). Angle bound is a setting only available in SLANTM. If angle bound is true, we bound the angle of rotation by a learnable magnitude value.